arXiv:1601.04213v2 [cs.DS] 25 Jan 2016

Partial-Match Queries with Random Wildcards: In Tries and Distributed Hash Tables

Junichiro Fukuyama

Applied Research Laboratory, The Pennsylvania State University
jxf140@psu.edu

Abstract. Consider an m-bit query q to a bitwise trie T. A *wildcard* $*$ is an unspecified bit in q for which the query asks the membership for both cases $* = 0$ and $* = 1$. It is common that such *partial-match queries* with wildcards are issued in tries. With uniformly random occurrences of w wildcards in q assumed, the obvious upper bound on the average number of traversal steps in T is $2^w m$. We show that the average does not exceed

$$\frac{m+1}{w+1}\left(2^{w+2} - 2w - 4\right) + m = O\left(\frac{2^w m}{w}\right),$$

and equals the value exactly when T includes all the m-bit keys as the worst case. Here the query q performs with the naive backtracking algorithm in T. It is similarly shown that the average is $O\left(\frac{k^w m}{w}\right)$ in a general trie of maximum out-degree k. Our analysis for tries is extended to a distributed hash table (DHT), which is among the most frequently used decentralized data structures in networking. We show, under a natural probabilistic assumption for the largest class of DHTs, that the average number of hops required by an m-bit query q to a DHT D with random w wildcards meets the same asymptotic bound. As a result, q is answered with average $O\left(\frac{2^w m}{w}\right)$ hops rather than $\Theta\left(2^w m\right)$ in the four major DHTs Chord, Pastry, Tapestry and Kademlia. In addition, with a uniform key distribution for sufficiently many entries, we prove that a lookup request to the DHT Chord is answered correctly with $O\left(m\right)$ hops and probability $1 - 2^{-\Omega(m)}$. To the author's knowledge, the probability $1 - 2^{-\Omega(m)}$ of correct lookup in Chord has not been identified so far.

Keywords: partial-match query, trie, distributed hash table, Chord, Kademlia, Tapestry, Pastry, Koorde, wildcard matching

1 Introduction

Finding information that partially matches to a given pattern has been a major problem in computer science for decades. In addition to the classical RK and KMP-algorithms in textbooks such as [1], a collection of research results on *partial-match queries* is found in literature such as [2,3]. It is common in practice to construct a trie as the data structure for partial-match queries with wildcards

[4]. Here a *trie* is the well-known prefix tree data structure to store keys [5], used for applications including dictionary search and lexicographic sorting. The most basic form of a trie T is the *bitwise trie* to store m-bit integer keys. Denote by q a query to such T. A *wildcard* $*$ in q is defined as an unspecified bit for which q asks the membership for both cases $*=0$ and $*=1$. For example, $q=1*0*0$ is a 5-bit query asking if 10000, 10010, 11000 and 11010 are in T.

In the paper, we first analyze the average performance of an m-bit query q to T with random w wildcards. Assume that T is a bitwise trie for which the wildcards occur in w positions in q, chosen randomly with the uniform probability density function (PDF), and also that q performs with the naive backtracking algorithm. We show that the average number of steps in T required by q does not exceed

$$\frac{m+1}{w+1}\left(2^{w+2}-2w-4\right)+m=O\left(\frac{2^w m}{w}\right),$$

and is exactly equal to the value when T includes all the m-bit keys as the worst case. This improves the obvious upper bound $2^w m$ asymptotically. We will also prove that the average is $O\left(\frac{k^w m}{w}\right)$ in a general trie of maximum out-degree k. The results have been unknown so far despite the common use of queries with wildcards in T. In Section 4, we will present an example of a practical system in which the above analysis could be useful.

The second half of the paper extends our analysis to a distributed hash table (DHT), which is among the most significant decentralized data structures used in networking. A DHT can support a number of application services such as web caching, file sharing, name-address mapping to track node mobility [6], instant messaging, multicast, content distribution, etc. In [7], detailed analysis is presented on the tradeoff between the routing table size and average number of hops per lookup (*network diameter*) in different DHTs. In the taxonomy, the class of DHTs with $O(\log n)$ routing table size and $O(\log n)$ network diameter is the largest one (n: the number of nodes in the DHT). We focus on this DHT class denoted by $\mathcal{C}$, which includes the four major DHTs Chord [10], Pastry [11], Tapestry [12], and Kademlia [13].

We will see a structural similarity between a bitwise trie and DHT in order to answer an m-bit query. With the above $O\left(\frac{2^w m}{w}\right)$ bound for bitwise tries and another probabilistic assumption, we show that the average number of hops required by an m-bit query q to a DHT D with random w wildcards meets the same asymptotic bound. Arguing that the probabilistic assumption holds generally for the DHT class $\mathcal{C}$, we will especially confirm it for the above four DHTs. The result thus improves the theoretical upper bound on the lookup time with w random wildcards from $O(2^w \log n)$ to $O\left(\frac{2^w \log n}{w}\right)$ in the four DHTs.

In addition, with a uniform key distribution for $\Omega(mn)$ entries, we prove that a lookup request to the DHT Chord is answered correctly with $O(m)$ hops and probability $1-2^{-\Omega(m)}$. The probability $1-2^{-\Omega(m)}$ of correct lookup in Chord will be identified for the first time to the author's knowledge.

The rest of the paper is structured as follows. In Section 2, we will prove the $O\left(\frac{2^w m}{w}\right)$ and $O\left(\frac{k^w m}{w}\right)$ bounds for tries T. Section 3 shows the $O\left(\frac{2^w m}{w}\right)$ bound

for the four DHTs, and the probability $1 - 2^{-\Omega(m)}$ of correct lookup in Chord. It is followed by concluding remarks in Section 4.

2 Average Search Time with Random Wildcards in Tries

2.1 In a Bitwise Trie

An *m-bit query q to a bitwise trie T with w wildcards* is a string consisting of $m - w$ 0s and/or 1s, and w wildcards *. We assume that the letters in q are numbered m, $m-1$, ..., 1 from the left to right (*bit positions*).

We measure the running time of a query by the number of edges in T traversed by the search algorithm, calling them *steps*. A query q with w wildcards completes in no more than $2^w m$ steps. We use the standard O, Ω and Θ-notations to express asymptotic quantities. A *constant* in this paper means a fixed positive real number depending on no other variable.

In this section, we prove that q with w random wildcards takes average $O\left(\frac{2^w m}{w}\right)$ steps. By w random wildcards, we mean the following uniform assumption.

Assumption I: In an m-bit query q with w wildcards, $*$ occurs in w positions with the uniform PDF.

In other words, every wildcard pattern, or *configuration*, occurs with the same probability $1/\binom{m}{w}$. Here a configuration determines the wildcard positions of a query q to T. For example, c*cc* is a configuration in which c represents 0/1. If q satisfies Assumption I, it is said to be a *query q to T with uniformly random w wildcards*.

Also consider the following natural backtracking search algorithm in T:

Algorithm QUERY**:** *Started at the root of T, search for the key such that every wildcard $*$ is 0 in q. When the current key's membership is determined, backtrack to the node of T representing the closest unfinished wildcard*[1] *$*$ in q. Change $* = 0$ into $* = 1$. Search for the new key in T. Continue until the memberships of all the 2^w keys are determined.*

The intuition behind the proof of $O\left(\frac{2^w m}{w}\right)$ steps is the following. If the wildcards in q occur in bit positions bounded by a small integer $j \geq 1$, it takes at most $m + O(2^w j)$ steps to answer q, which is much smaller than $2^w m$. Since wildcards are placed randomly with the uniform PDF, this must affect the asymptotic number of steps required by q.

We prove the proposition below. It will be extended to a general trie of maximum out-degree k in the next subsection.

Proposition 1. *Let q be an m-bit query to a trie T with uniformly random w wildcards. Algorithm* QUERY *answers q in no more than*

$$b = \frac{m+1}{w+1}\left(2^{w+2} - 2w - 4\right) + m$$

[1] This means $*$ such that the search for $* = 0$ is finished but $* = 1$ is not.

average steps, and in exactly b average steps when T includes all the possible m-bit keys. □

We construct its proof in what follows. We first show that the average is at most

$$s(m,w) \stackrel{def}{=} m + \sum_{\substack{1\le z\le m\\ 1\le j\le w}} j2^{w-j+1}\frac{\binom{z}{j}\binom{m-z}{w-j}}{\binom{m}{w}}. \tag{1}$$

The fraction $\binom{z}{j}\binom{m-z}{w-j}\Big/\binom{m}{w}$ is well-known as the *hypergeometric distribution* [8]. It is the probability of j successes in w draws without replacement, from m items including z successes and $m-z$ failures.

Let z_1 denote the position of the rightmost wildcard bit, *i.e.*, the *least significant wildcard* in the given query q. Likewise, let z_j be the position of the j^{th} least significant wildcard. The set

$$Z = \{z_1, z_2, \ldots, z_w\}$$

determines a configuration of q. Observe a lemma first for q with a fixed configuration.

Lemma 1. *Algorithm* QUERY *on a query q with configuration* $\{z_1, z_2, \ldots, z_w\}$ *terminates in*

$$\hat{s}\,(m,w) \stackrel{def}{=} m + \sum_{j=1}^{w} 2^{w-j+1} z_j \tag{2}$$

steps or less.

Proof. Prove the claim by induction on w. The basis occurs when $w = 1$. One can check that q with one wildcard in position z_1 takes at most $m+2z_1 = \hat{s}\,(m,1)$ steps, verifying the basis.

Assume true for $w-1$ and prove true for w. Below the stated number of steps are all in the worst case. The algorithm QUERY first sets the most significant wildcard at z_w as $* = 0$, and performs the m-bit query with $w-1$ wildcards. It takes $\hat{s}\,(m, w-1)$ steps by induction hypothesis. Then it backtracks to the node representing the bit position z_w with z_w steps, set $* = 1$, and recursively search for the remaining $w-1$ wildcards again. This takes extra $\hat{s}\,(z_w, w-1)$ steps.

So the total number of steps required by q is at most

$$\begin{aligned}
&\hat{s}\,(m, w-1) + z_w + \hat{s}\,(z_w, w-1)\\
&= \left(m + \sum_{j=1}^{w-1} 2^{w-j} z_j\right) + z_w + \left(z_w + \sum_{j=1}^{w-1} 2^{w-j} z_j\right)\\
&= m + \sum_{j=1}^{w} 2^{w-j+1} z_j = \hat{s}\,(m,w)\,,
\end{aligned}$$

proving the induction step. The lemma follows. □

Next we calculate the average of $\hat{s}\,(m, w)$ with the uniform occurrence of $Z = \{z_1, z_2, \ldots, z_w\}$. Let $z \le m$ and $j \le w$ be positive integers. Denote by $p_{z,j}$ the probability that z is the position of the j^{th} least significant wildcard. We have the following lemma.

Lemma 2. $p_{z,j} = \binom{z-1}{j-1}\binom{m-z}{w-j}\Big/\binom{m}{w}$.

Proof. Fix z and j. The number of Z such that $z_j = z$ is $\binom{z-1}{j-1}\binom{m-z}{w-j}$. Since each configuration occurs with probability $1/\binom{m}{w}$, the probability of $z_j = z$ is $p_{z,j}$ as claimed. □

If a given integer z is z_j in (2), it causes $2^{w-j+1}z_j = 2^{w-j+1}z$ steps in the summation, which occurs with the probability $p_{z,j}$. The average number of steps required by q is thus bounded by

$$m + \sum_{\substack{1\le z\le m\\ 1\le j\le w}} 2^{w-j+1} z p_{z,j} = m + \sum_{\substack{1\le z\le m\\ 1\le j\le w}} 2^{w-j+1} z \binom{z-1}{j-1}\binom{m-z}{w-j}\Big/\binom{m}{w}$$

$$= m + \sum_{\substack{1\le z\le m\\ 1\le j\le w}} j 2^{w-j+1} \binom{z}{j}\binom{m-z}{w-j}\Big/\binom{m}{w} = s(m, w).$$

This proves our claim that the algorithm QUERY takes at most $s(m, w)$ steps on average. The bound is tight; when T includes all the possible m-bit keys, QUERY actually takes average $s(m, w)$ steps.

It now suffices to show

$$s(m, w) = \frac{m+1}{w+1}\left(2^{w+2} - 2w - 4\right) + m, \tag{3}$$

to prove Proposition 1. As in standard textbooks on generating functions such as [9], for a function $f : (0, 1) \to \mathbb{R}$ with its Taylor series, denote by $[x^k]f(x)$ the coefficient of x^k in the series. Since $\binom{z}{j} = [x^z]\frac{x^j}{(1-x)^{j+1}}$ and $\binom{m-z}{w-j} = [x^{m-z}]\frac{x^{w-j}}{(1-x)^{w-j+1}}$,

$$\sum_{1\le z\le m} \binom{z}{j}\binom{m-z}{m-j} = \sum_{0\le z\le m} \binom{z}{j}\binom{m-z}{m-j}$$

$$= [x^m]\frac{x^j}{(1-x)^{j+1}} \cdot \frac{x^{w-j}}{(1-x)^{w-j+1}} = [x^m]\frac{x^w}{(1-x)^{w+2}}$$

$$= [x^{m+1}]\frac{x^{w+1}}{(1-x)^{w+2}} = \binom{m+1}{w+1}.$$

So,

$$s(m, w) = m + \sum_{1\le j\le w} j 2^{w-j+1}\frac{\binom{m+1}{w+1}}{\binom{m}{w}} = \frac{m+1}{w+1}\left(2^{w+2} - 2w - 4\right) + m,$$

verifying (3). As we have already confirmed that the bound is tight, this completes the proof of Proposition 1.

2.2 In a Trie of Maximum Out-Degree k

We now consider a general trie T of maximum out-degree k. We generalize Proposition 1 into:

Theorem 1. *Let T be a trie with maximum out-degree $k \geq 2$. A query to T of length m with w uniformly random wildcards can be answered in average*

$$b = m + \frac{2(m+1)}{w+1} \cdot \frac{k^{w+1} - (w+1)k + w}{k-1}$$

steps or less. The average is exactly b when q performs with Algorithm Query*, and T is a complete k-ary tree.* □

This means q requires $O\left(\frac{k^w m}{w}\right)$ steps in T as claimed in the introduction.

A general trie T is formally defined with its *membership*: It is a tree such that each edge is associated with a letter in a given set A (*alphabet*). A string $s \in A^*$ is said to be a *member of* T if there exists a maximal directed path $\{e_1, e_2, \ldots, e_n\}$ in T such that s is the concatenation of the letters given on the edges $e_1, e_2, \ldots, e_n$ in the order. For such T, the algorithm Query is naturally generalized. We re-define $\hat{s}(m, w)$ in (2) by

$$\hat{s}\,(m, w) = m + \sum_{j=1}^{w} 2k^{w-j}(k-1)z_j.$$

We show the same claim as Lemma 1 with the new $\hat{s}(m, w)$.

Lemma 3. *Algorithm* QUERY *on a given query q having a configuration $Z = \{z_1, z_2, \ldots, z_w\}$ takes no more than $\hat{s}\,(m, w)$ steps.*

Proof. Prove by induction on w. The basis $w = 1$ is straightforward to check. Assume true for $w-1$ and prove true for w. It suffices show that the number of steps required by q is at most $\hat{s}\,(m, w-1) + (k-1)z_w + (k-1)\hat{s}\,(z_w, w-1)$ since

$$\begin{aligned}
&\hat{s}\,(m, w-1) + (k-1)z_w + (k-1)\hat{s}\,(z_w, w-1) \\
&= \left(m + \sum_{j=1}^{w-1} 2k^{w-j-1}(k-1)z_j\right) + (k-1)z_w + (k-1)\left(z_w + \sum_{j=1}^{w-1} 2k^{w-j-1}(k-1)z_j\right) \\
&= m + 2(k-1)z_w + k\sum_{j=1}^{w-1} 2k^{w-j-1}(k-1)z_j \\
&= m + \sum_{j=1}^{w} 2k^{w-j}(k-1)z_j = \hat{s}\,(m, w)\,.
\end{aligned}$$

To verify it, wlog let v be the node such that the letters $a_1, a_2, \ldots, a_k$ given on the edges from v correspond to the most significant wildcard in q. The algorithm QUERY first chooses a_1 as the value of the wildcard, then finds all the members of T matching to q. This requires at most $\hat{s}(m, w-1)$ steps by induction hypothesis. Then it backtracks to v in z_w steps to find all the members of T that match to q including a_2. It takes at most $z_w + \hat{s}(z_w, w-1)$ steps.

The above repeats $k-1$ times for $a_2, a_3, \ldots, a_k$. Thus the total number of traversal steps required by q is upper-bounded by $\hat{s}\,(m, w-1) + (k-1)z_w + (k-1)\hat{s}\,(z_w, w-1)$, completing the proof. □

The rest of the proof is the same as for a bitwise trie. We find that the average number of steps required by q is no more than

$$\begin{aligned}
&m + \sum_{\substack{1 \le z \le m \\ 1 \le j \le w}} 2k^{w-j}(k-1)zp_{z,j} = m + \sum_{\substack{1 \le z \le m \\ 1 \le j \le w}} 2jk^{w-j}(k-1)\frac{\binom{z}{j}\binom{m-z}{w-j}}{\binom{m}{w}} \\
&= m + \sum_{1 \le j \le w} 2jk^{w-j}(k-1)\frac{\binom{m+1}{w+1}}{\binom{m}{w}} \\
&= m + \frac{2(m+1)}{w+1} \cdot \frac{k^{w+1} - (w+1)k + w}{k-1}.
\end{aligned}$$

The bound is tight by the same argument also; for q having a configuration Z, *Query* takes $\hat{s}\,(m, w)$ steps exactly if T is a complete k-ary tree. This completes the proof of Theorem 1.

3 Lookup Response Time with Wildcards in a Distributed Hash Table

In this section, we show the same asymptotic upper bound $O\left(\frac{2^w m}{w}\right)$ for a DHT D. We will verify it through the structural similarity between a bitwise trie

and DHT: key search by incremental bit improvement. We first define general terminology on a DHT with related facts in Section 3.1. The second subsection presents a necessary probabilistic assumption general in the aforementioned DHT class $\mathcal{C}$. In Section 3.3, we show that the probability of correct lookup is $1 - 2^{-\Omega(m)}$ in the DHT Chord with sufficiently many independent keys. The $O\left(\frac{2^w m}{w}\right)$ bound will be proved with Proposition 1 in Section 3.4.

3.1 Distributed Hash Table and Wildcard Query

Let S be the *key space* for DHT D. Suppose it consists of the m-bit binary integers so that $|S| = 2^m$. A node v in D is labeled by a key denoted by $key(v) \in S$. It is said to be the *node key* of v, which is typically a large random number such as a hash of the IP address of v or that of a file name. The mapping $v \mapsto key(v)$ is an injection, *i.e.*, there is no other node v' in D such that $key(v') = key(v)$. Information is stored at a node as a pair <key, value> called *entry*. We denote an entry by $<d, r>$ where $d \in S$ is its *data key*.

The *distance from $d \in S$ to $d' \in S$* is written as $\Delta(d, d')$, which is defined by the DHT design. For example, Chord measures $\Delta(d, d')$ as $d' - d \bmod 2^m$ evaluated clockwise in the circular ring $0, 1, \ldots, 2^{m-1}$ [10]. Kademlia measures $\Delta(d, d')$ by the XOR metric [13]. For a data key d, we say that the node v such that $\Delta(d, key(v))$ is minimum is the *successor of d*, and v such that $\Delta(key(v), d)$ is minimum is the *predecessor*. An entry $<d, r>$ is stored at the successor or predecessor of d, or in a generalized object to include them. Also the *successor of v* is the node $v' \neq v$ such that $\Delta(key(v), key(v'))$ is minimum, and *predecessor of v* is $v' \neq v$ such that $\Delta(key(v'), key(v))$ is minimum. A *neighbor* of d or v is its successor or predecessor. Denote by n the number of nodes in D. We assume $m = O(\log n)$ conventionally.

It is called *lookup in D* to determine the membership of a given key $d \in S$ in D, written as $d \in D$ or $d \notin D$. To answer it, the lookup protocol runs at the *current peer node* of D moving to another if necessary. A *hop* is a change of the current peer node. The average number of hops per lookup is said to be the *network diameter of D*.

In addition, each node v holds a set of addresses of other nodes determined by certain rules, usually including all the v's neighbors. It is called the *routing table of v*. We may simply say the routing table includes the nodes rather than their addresses. A good DHT is designed with a routing table and distance $\Delta(d, d')$ that allow for efficient lookups and updates of entries. For D in the aforementioned DHT class $\mathcal{C}$, the routing table size and network diameter are both $O(\log n)$. The class $\mathcal{C}$ includes the four major DHTs Chord, Pastry, Tapestry and Kademlia.

With the above, a *query q to D with uniformly random w wildcards* is defined the same way as to a bitwise trie T. We say that a lookup/query is *resolved* if the protocol returns the correct answer. Let h stand for the average number of hops required to resolve q. It is our measure of q's response time. Our goal in this section is to show $h = O\left(\frac{2^w \log n}{w}\right)$ for DHT $D \in \mathcal{C}$ and q with uniformly random w wildcards.

3.2 A Probabilistic Assumption for D

It works similarly to a bitwise trie T how to find a data key d that is a member of DHT D: by repeatedly moving to a node v such that $key(v)$ has a smaller distance to d than the current peer node. The number of significant bits shared by d and $key(v)$ is increased incrementally. In our proof of the $O\left(\frac{2^w m}{w}\right)$ bound for D, we need another probabilistic assumption to justify this incremental bit improvement.

A DHT D or its lookup protocol is said to *improve at least one bit per hop, correctly with high probability* if it satisfies the three conditions A)–C) below: In finding a target data key d, let v be the current peer node and v_d be a neighbor of d. Suppose that $2^{g-1} \leq \Delta\left(key(v), key(v_d)\right) < 2^g$ for an integer $g > 1$. Let v' be a node in the routing table such that

$$2^{g-2} \leq \Delta\left(key(v'), key(v_d)\right) < 2^{g-1}. \tag{4}$$

The three conditions are:

A) The routing table of v include v' such that (4) with probability $1 - 2^{-m^{1/2+\epsilon}}$ for some sufficiently small constant $\epsilon > 0$.
B) If there exists such v', the lookup protocol must move the current peer node to v'.
C) The worst case number of hops for the lookup does not exceed a polynomial in m.

If the routing table of v does not include such v', the lookup protocol may decide $d \notin D$, or change v to another with no guarantee on the closeness to d. This *error case* occurs with a small probability at most $2^{-m^{1/2+\epsilon}}$ for each v.

Our assumption for the proof of $h = O\left(\frac{2^w \log n}{w}\right)$ is now stated as:

Assumption II: The lookup protocol of the considered DHT improves at least one bit per hop, correctly with high probability.

This property of incremental bit improvement is common in the considered DHT class $\mathcal{C}$. The lookup protocol keeps improving another bit until d is between the node keys of v and its neighbor for the first time. In the end it identifies both the successor and predecessor of d. An error case may occur with probability $p(m)2^{-m^{1/2+\epsilon}} < 2^{-m^{(1+\epsilon)/2}}$ for some polynomial $p(m)$. Thus any lookup in D satisfying the assumption is resolved with average number of hops $O(m) = O(\log n)$, and probability at least $2^{-m^{(1+\epsilon)/2}}$.

Hence Assumption II is general in $\mathcal{C}$, and is satisfied by the above four DHTs: One can check that all of their lookup protocols improve at least one bit per hop with high probability. The actual magnitude of the high probability depends on m, n, and the frequencies of entry updates and routing table maintenance. Assumption II with the bound $1 - 2^{-m^{1/2+\epsilon}}$ is true for the four DHTs with some possible performance parameters in practice. Notice that if $m = C \log n$ for a constant C, it means $1 - 2^{-m^{1/2+\epsilon}} \leq 1 - n^{-\epsilon}$ for any small constant $\epsilon > 0$ and

sufficiently large m and n. The error probability bound $n^{-\epsilon}$ can be achieved in any of the four DHTs. Also the condition C) is satisfied by the maximum number of hops allowed for a lookup, which is set in the DHT.

It has been seen that D satisfying Assumption II searches for keys with the same incremental bit improvement as a bitwise trie T. Hence we will be able to apply Proposition 1 to D to show the $O\left(\frac{2^w m}{w}\right)$ bound. Here the following *natural query protocol* is assumed for D, which is equivalent to the algorithm QUERY.

Natural Query Protocol: *First set every $*$ in q as 0 and search for the data key in the DHT. Change the least significant unfinished wildcard from $* = 0$ into $* = 1$. Search for the new data key started at the current peer node. Repeat until the membership of every desired data key is determined.*

Note that we consider two independent probability spaces for a) the key distribution in D, and b) the distribution of configurations of q. If we say the average number of hops for q in D, it means the average over the joint distribution decided by a) and b).

3.3 The Probability of Correct Lookup in the DHT Chord

In case D is Chord, we can present a parameter class such that D satisfies Assumption II exactly. Consider the following argument.

Lemma 4. *Let D be a distributed hash table Chord defined over the m-bit key space S with n nodes where n and $m = O(\log n)$ are sufficiently large. D satisfies Assumption II if:*

i) there are at least Cmn entries stored in D for a sufficiently large constant $C > 1$, and

ii) an entry $<d, r>$ is stored at a node chosen with the uniform probability density function, independently of the others[2]*.*

Proof. It suffices to show that there are at least m entries stored at any given node v with high probability, which is seen as follows. By the construction of Chord [10], the i^{th} entry stored at v has a pointer to the successor of $key(v)+2^{i-1}$, called *finger*. In other words, v's routing table is required to include the address of the successor if there are i entries or more stored at v. If there are m entries at v with high probability, its routing table has the finger to the successor of $key(v) + 2^{i-1}$ for every $i \leq m$. Then the lookup protocol defined by Chord improves at least one bit per hop correctly with high probability[3].

[2] This statement considers a probability space constructed for each given m, n and the number of stored entries. Its event set consists of all the cases of contained node keys and entries. It defines a PDF of node choice to store each entry. It is uniform and independent of any other event, as the statement assumes.

[3] If the routing table of v includes no other node v' closer to the desired key d (*i.e.*, such that (4)), the protocol of Chord decides $d \notin D$, rather than performing further lookup with no guarantee to the closeness to d.

Let N be the total number of entries in D that is at least Cmn by Condition i), and j be the number of entries stored at v. Due to ii), deciding if the i^{th} entry is stored at v is a Bernoulli trial with probability of success equal to $\frac{1}{n}$. Repeating it N times, we have $Pr(j \leq m) = \sum_{j \leq m} \binom{N}{j} \left(\frac{1}{n}\right)^j \left(1 - \frac{1}{n}\right)^{N-j}$, where $Pr(\cdot)$ denotes the probability of the argument event. We will show

$$Pr\left(j \leq m\right) < e^{-\frac{C}{2}m}. \tag{5}$$

Then $Pr\left(j \leq m\right) < e^{-\frac{C}{2}m} < e^{-m^{1/2+\epsilon}}$, meaning v has m entries with high probability as required by Assumption II. (Note that the assumption considers a single particular hop from the current peer node v.)

We show (5) by the Chernoff bound given in [15]. For our case, it provides the upper bound

$$Pr\big(X_1 + X_2 + \cdots + X_N \leq m\big) \leq \min_{t \leq 0} \left\{ e^{-tm + N \ln M(t)} \right\}, \tag{6}$$

$$\text{where } M(t) = \left(1 - \frac{1}{n}\right) e^{0 \cdot t} + \frac{1}{n} e^{1 \cdot t} = 1 + \frac{e^t - 1}{n}.$$

Here X_i is the random variable that represents the i^{th} Bernoulli trial, *i.e.*, $X_i = 1$ if i^{th} entry is stored at v and $X_i = 0$ otherwise. Also $M(t)$ is the *moment generating function of* X_i where t is a real parameter.

By (6), $Pr\left(j \leq m\right) \leq e^{-tm + N \ln M(t)}$ for the parameter $t = \ln \frac{mn}{N} \leq \ln \frac{1}{C} < 0$ that is particularly chosen. The natural logarithm of the moment generating function is $\ln M(t) = \ln \left(1 + \frac{\frac{mn}{N} - 1}{n}\right)$ for this t. We now have

$$\ln Pr\left(j \leq m\right) < -m \ln \frac{mn}{N} + N \ln \left(1 + \frac{\frac{mn}{N} - 1}{n}\right), \tag{7}$$

desiring that its RHS is at most $-Cm/2$ to show (5).

Put $y = \frac{N}{mn} \geq C$ that is sufficiently large. By the Taylor series of the natural logarithm, $\ln \left(1 + \frac{\frac{mn}{N} - 1}{n}\right) = \ln \left(1 - \frac{1 - \frac{1}{y}}{n}\right) \leq -\frac{1 - \frac{1}{y}}{n} + O\left(\frac{1}{n^2}\right)$. By (7) and $N = mny$,

$$\begin{aligned} \ln Pr\left(j \leq m\right) &< m \ln y - \frac{N}{n}\left(1 - \frac{1}{y}\right) + O\left(\frac{N}{n^2}\right) \\ &= m\left(\ln y - y\left(1 - \frac{1}{y}\right) + O\left(\frac{y}{n}\right)\right) \\ &= m\left(\ln y - y + 1 + O\left(\frac{y}{n}\right)\right) < -\frac{y}{2}m \leq -\frac{C}{2}m. \end{aligned}$$

This confirms (5) proving the lemma. □

Observe that if Chord D satisfies Conditions i) and ii), an error case occurs for each lookup with probability at most $m \cdot e^{-\frac{C}{2}m} < 2^{-\Omega(m)}$ due to (5). Then any lookup in D is resolved correctly with at most $m = O\left(\log n\right)$ hops and probability $1 - 2^{-\Omega(m)}$. Therefore:

Theorem 2. *Let D be a distributed hash table Chord defined over the m-bit key space S with n nodes where n and $m = O(\log n)$ are sufficiently large. Suppose that it satisfies the following two.*

i) There are at least Cmn entries stored in D for a large constant $C > 1$.

ii) An entry $<d, r>$ is stored at a node chosen with the uniform probability density function, independently of the others.

Then any lookup in D is resolved with at most $m = O(\log n)$ hops and probability $1 - 2^{-\Omega(m)}$. □

The theorem confirms the aforementioned bound $1-2^{-\Omega(m)}$. In other words, the sufficient condition for a successful lookup in Chord with the probability bound is i) and ii), which assumes that there are enough entries in D created by a series of mutually independent N Bernoulli trials.

3.4 Proof of the $O\left(\frac{2^w \log n}{w}\right)$ Bound

We now show our main claim.

Theorem 3. *Let D be a distributed hash table defined over the m-bit key space, which improves at least one bit per hop correctly with high probability, and let q be a query to D with uniformly random w wildcards. Then the natural query protocol resolves q with high probability, and with the average number of hops at most $O\left(\frac{2^w m}{w}\right)$.*

Proof. Denote by h the average number of hops, and by $d_1, d_2, \ldots, d_{2^w}$ the 2^w data keys specified by q in the order determined by the natural query protocol. We first show

$$h \leq \frac{m+1}{w+1}\left(2^{w+2} - 2w - 4\right) + m + O\left(2^{-m^{(1+\epsilon)/2}}\right). \tag{8}$$

Observe facts on h and d_i.

a) At most m hops are necessary to determine if $d_1 \in D$, and j hops to determine if $d_i \in D$ for $i > 1$, where j is the position of the unfinished least significant wildcard in q when the lookup for d_{i-1} is complete. By Assumption II, this is true except for an error case occurring with probability $2^{m^{1/2+\epsilon}}$ or less.

b) In an error case, the total number of hops is bounded by 2^w times a polynomial in m. Its contribution to h is the $O\left(2^{-m^{(1+\epsilon)/2}}\right)$ term in (8). We ignore it in the arguments below.

c) Denote by h_i the number of extra hops required for d_i considered in a). To compare it with traversal steps in a bitwise trie T, let s_i be the worst case number of extra steps necessary for Algorithm QUERY to determine if d_i is in T, after the search for d_{i-1} is complete. We have

$$h_i \leq s_i \text{ for } i = 1, 2, \ldots, 2^w :$$

If $i = 1$ then $h_1 = s_1 = m$, otherwise $h_i = j$ and $s_i = 2j$ where j is the same as in a).

d) Let b be as given by Proposition 1. It upper-bounds the average number of steps required by q in T. Thus $\mathbb{E}\left(\sum_{i=1}^{2^w} s_i\right) \leq b$, where $\mathbb{E}(\cdot)$ denotes the average of the argument random variable.

Hence we have

$$h = \mathbb{E}\left(\sum_{i=1}^{2^w} h_i\right) \leq \mathbb{E}\left(\sum_{i=1}^{2^w} s_i\right) \leq b = \frac{m+1}{w+1}\left(2^{w+2} - 2w - 4\right) + m,$$

proving (8).

It remains show that the natural query protocol resolves q with high probability. If no further bit is improved at the current peer node v, the protocol may decide that $d_i \notin D$ or change v to another with no closeness guarantee. Such an error case occurs with probability at most $2^{-m^{1/2+\epsilon}}$ by Assumption II. The total number of hops is at most 2^w times a polynomial in m, say $p(m)$. An error case occurs at any peer node with probability no more than $2^{-m^{1/2+\epsilon}} \cdot 2^w p(m) < 2^{-m^{(1+\epsilon)/2}}$. Therefore, the protocol returns the correct answer to q with probability at least $1 - 2^{-m^{(1+\epsilon)/2}}$, a high probability. The theorem follows this statement. □

As stated in Section 3.1, we assume $m = O(\log n)$ in a DHT, so the theorem means $h = O\left(\frac{2^w \log n}{w}\right)$ as desired. The bound is applicable to Chord, Pastry, Tapestry and Kademlia since they satisfy Assumption II.

We note that the bound could also improve the performance of 2^w independent lookups in Koorde [16]: Koorde is a variant of Chord with the use of De Bruijin graph, achieving $O\left(\log n / \log\log n\right)$ hops per lookup with $O(\log n)$ routing table size. If 2^w lookups run independently in Koorde, its number of hops is $O\left(\frac{2^w \log n}{\log\log n}\right)$ whose argument is greater than $O\left(\frac{2^w \log n}{w}\right)$ when w is sufficiently larger than $\log\log n$.

4 Concluding Remarks and Open Problems

We have shown the bound $O\left(\frac{2^w m}{w}\right)$ for both bitwise tries and distributed hash tables in $\mathcal{C}$, and $O\left(\frac{k^w m}{w}\right)$ for a general trie of maximum out-degree k. They limit the asymptotic running time required by a partial-match query of length m with w uniformly random wildcards. We also confirmed the probability $1 - 2^{-\Omega(m)}$ of correct lookup in Chord under the natural assumption.

There are some practical cases to which the obtained results can be applied with the assumption of uniform wildcard occurrences. One such case is data retrieval: Suppose that one searches for data records with m attributes, managed as a trie T such that each attribute takes at most k values. An example of such a data record is of form `<college, department, building, title, last name, first name>`. The $m = 6$ attributes are hierarchical but with the

independent equal probability $\frac{1}{m}$ to be a wildcard in a query. The trie T organizing such data records could be an auxiliary data structure to enhance the search speed. In this situation, wildcards included in a query q occur randomly with the uniform PDF. By Theorem 1, q takes average $O\left(\frac{k^w m}{w}\right)$ steps rather than $k^w m$.

Further research on this problem could consider query protocols to resolve q with non-uniform probability distributions of wildcard occurrence. It is possible that such a protocol runs at multiple peer nodes simultaneously. It would be interesting to investigate its lookup efficiency.

Acknowledgements

The author is especially thankful to Yan Shvartzshnaider and Max Ott for introducing him to this problem. Also the author would like to thank Professor Hisashi Kobayashi at Princeton University, and Processors Martin Fürer and Piotr Berman at Penn State for their helpful suggestions.

References

1. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to Algorithms, Third Edition. MIT Press (2009).
2. Rivest, R.L.: Parttial-match retrieval algorithms. SIAM Journal on Computing, vol. 5, pp. 15-50 (1976).
3. Charikar, M., Indyk, P., Panigrahy, R.: New Algorithms for Subset Query, Partial Match, Orthogonal Range Searching, and Related Problems. LNCS, vol. 2380, pp. 451–463, Springer, Heidelberg (2002)
4. Sedgewick, R., Wayne, K.: Algorithms (4th edn.). Addison-Wesley (2011).
5. Knuth, D.E.: The art of computer programming, volume 3. Addison-Wesley (1997).
6. Vu, T., Baid, A., Zhang, Y., Nguyen, T. D., Fukuyama, J., Martin, R. P., Raychaudhuri, D.: DMap: A shared hosting scheme for dynamic identifier to locator mappings in the global internet. In: Proc. of the 32nd International Conference on Distributed Computing Systems (ICDCS12). IEEE, pp. 698–707 (2012).
7. Xu, J., Kumar, A., Yu, X.: On the fundamental tradeoffs between routing table size and network diameter in peer-to-peer networks. IEEE Journal on Selected Areas in Communications, vol. 22, pp. 151–163 (2004).
8. Rosenkrantz, W. A.: Introduction to probability and statistics for science, engineering, and finance. Chapman and Hall (2009).
9. Graham, R. L., Knuth, D. E., Patashnik, O.: Concrete mathematics. Addison-Wesley (1994).
10. Stoica, I., Morris, R., Karger, D., Kaashoek, M. F., Balakrishnan, H.: Chord: a scalable peer-to-peer lookup service for internet applications. In: Proc. of the SIGCOMM 2001 Conference (SIGCOMM01). ACM, vol. 31, pp. 149–160, New York (2001).
11. Rowstron, A., Druschel, P.: Pastry: scalable, decentralized object location and routing for large-scale peer-to-peer systems. In: Proc. of IFIP/ACM International Conference on Distributed Systems Platforms (Middleware). LNCS, vol. 2218, pp. 329–350. Springer, Heidelberg (2001).

12. Zhao, B.Y., Huang, L., Stribling, J., Rhea, S.C., Joseph, A.D., Kubiatowicz, J.D.: Tapestry: a resilient global-scale overlay for service deployment. IEEE Journal on Selected Areas in Communications vol. 22, pp. 41–53 (2004).
13. Maymounkov, P., Mazi'eres, D.: Kademlia: a peer-to-peer information system based on the XOR metric. In: Proc. of the 1st International Workshop on Peer-to Peer Systems (IPTPS02). LNCS, vol. 2429, pp. 53–65. Springer, Heidelberg (2002).
14. Graham, R. L., Knuth, D. E., Patashnik, O.: Concrete mathematics: a foundation for computer science (2nd edn.). Addison-Wesley (1994).
15. Kobayashi, H., Mark, L.M., Turin, W.: Probability, random processes, and statistical analysis: applications to communications, signal processing, queueing theory and mathematical finance. Cambridge University Press (2012).
16. Kaashoek, M.F., Karger, D.R.: Koorde: a simple degree-optimal distributed hash table. In: Proc. of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03). LNCS, vol. 2735, pp. 98–107. Springer, Heidelberg (2003).